Using IPMI Platform Management In Modular Computer Systems

Tom Slaight

Server Management Architect
Enterprise Platforms Group

Intel Corporation

Agenda

- IPMI in modular architecture
- IPMI specification support for modular
- IPMI components for modular systems
- IPMI future directions for modular
- Summary

Introduction

- **Audience:**
  Architects, Technical Managers, Firmware Leads, and Hardware Designers
  - **Involved in architecture, component selection, debug, test, or design of server baseboard and peripheral management subsystems**
- **Focus: IPMI-based platform management for modular system architectures**
  - **Hardware and Software components**

Early look at new IPMI Extensions

IPMI

Intelligent Platform Management Interface

- **Defines a standardized, abstracted, message-based interface to intelligent platform management hardware**
- **Defines standardized records for describing platform management devices and their characteristics**

Promoters: intel, HEWLETT PACKARD, NEC, DELL

Adopters: 145 and growing

http://developer.intel.com/design/servers/ipmi

intel

Intel Developer Forum Spring 2003

Initiative News

developer.intel.com/design/servers/ipmi

IPMI in modular architecture

Where it fits…

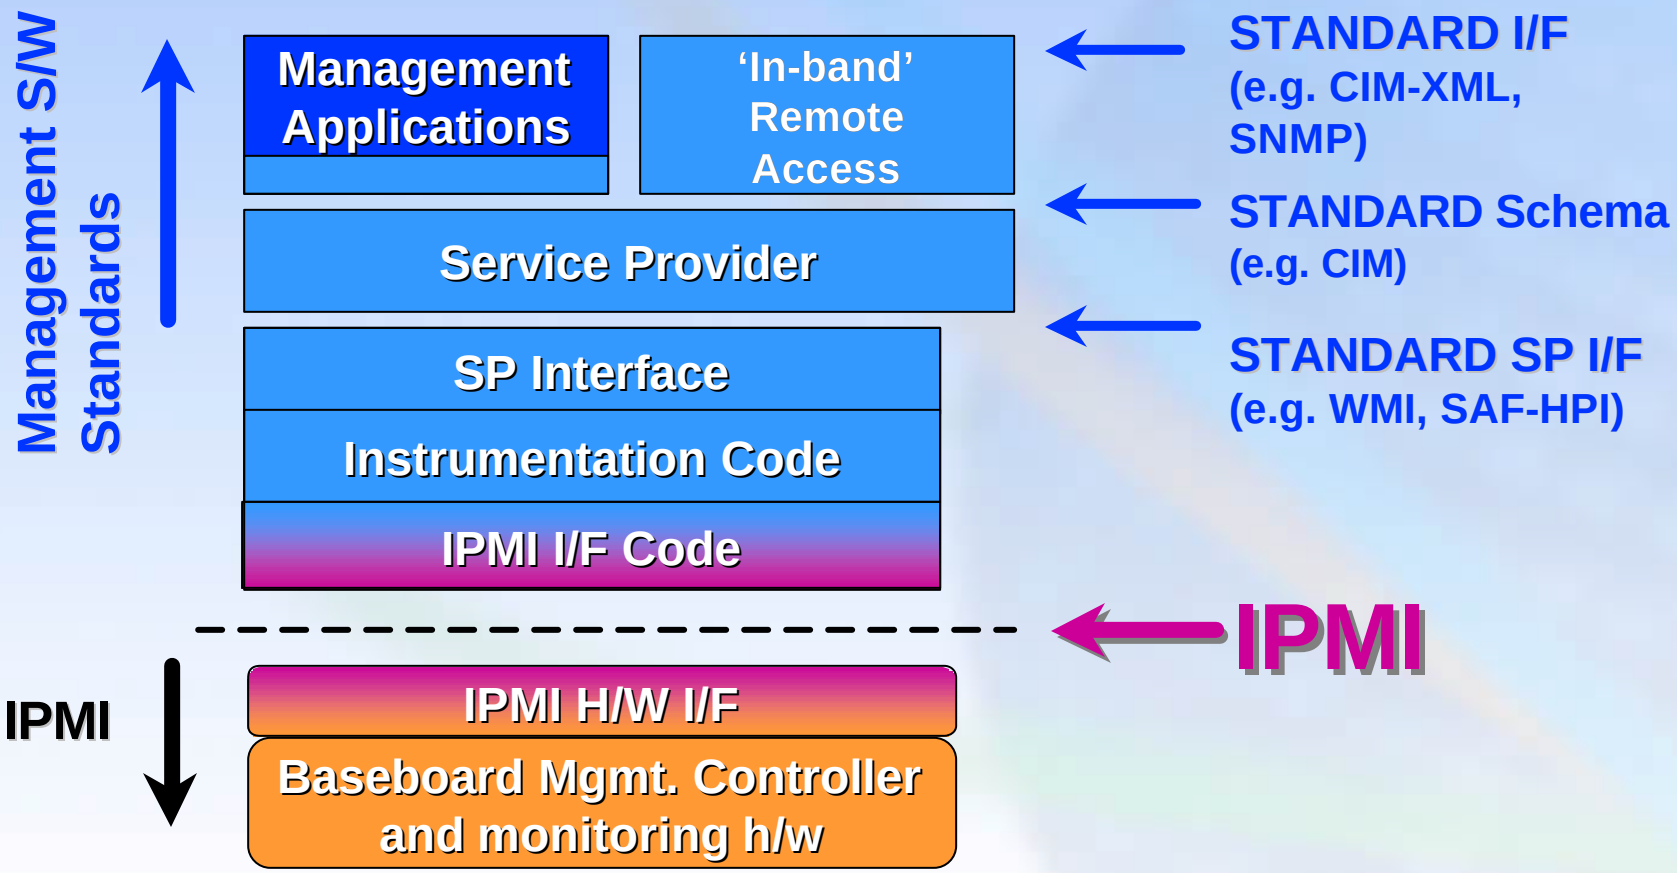

Complements existing management standards

IPMI in modular architecture

Where it fits…

The Modular Computing Data Center

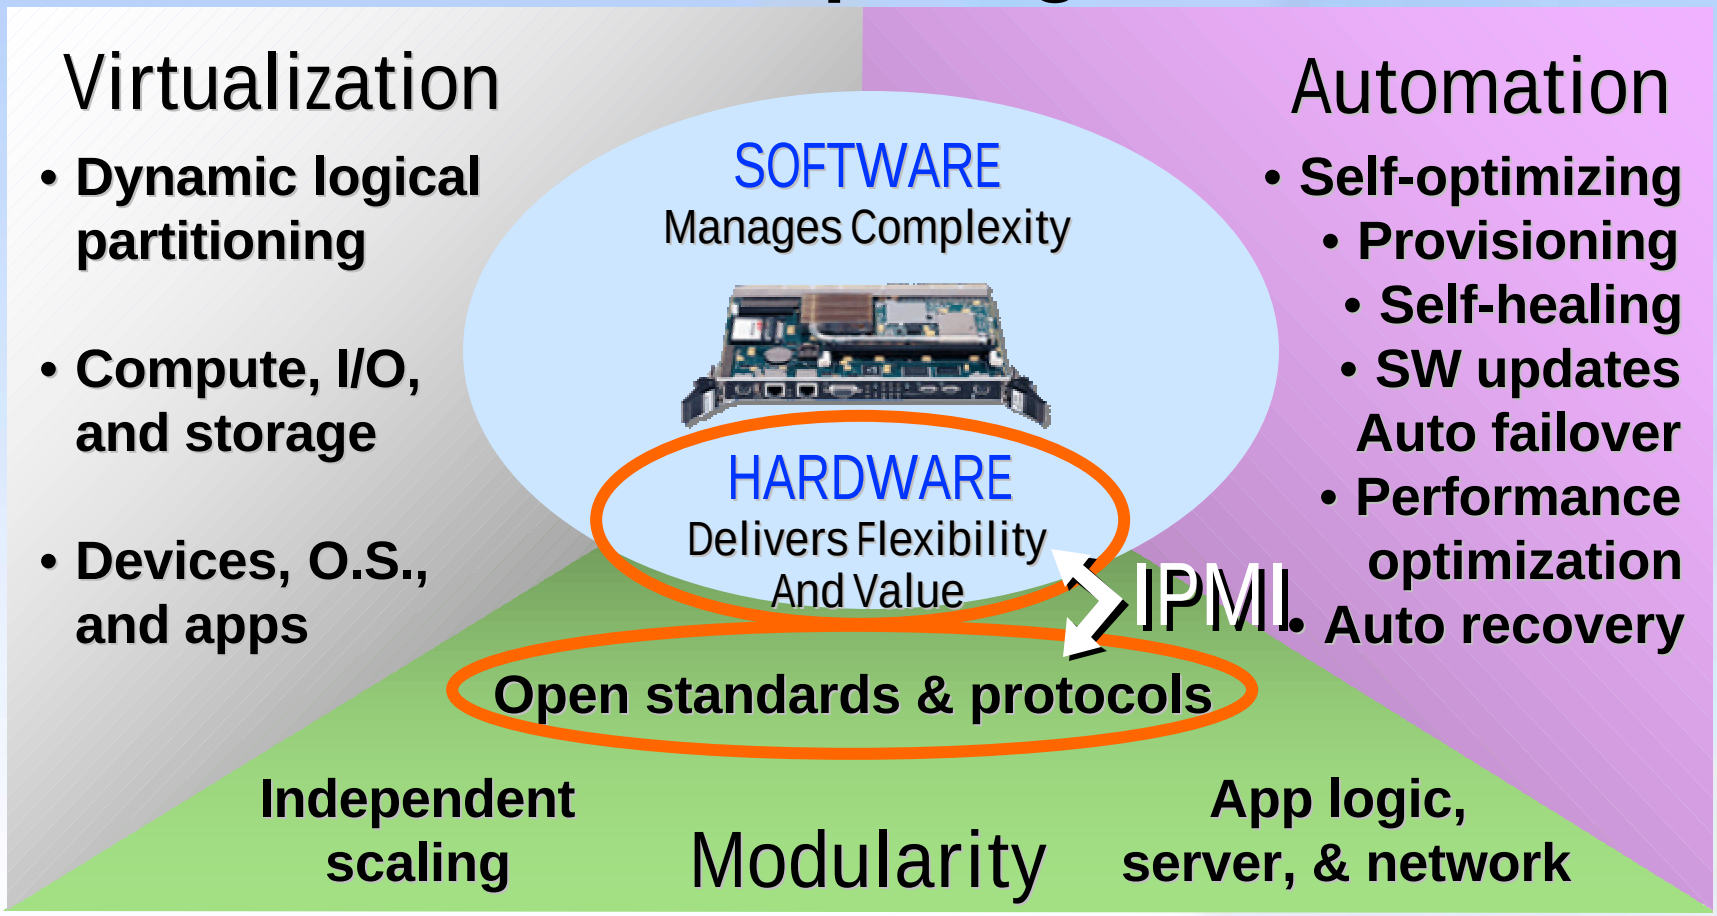

Enables Modular Hardware Platform

IPMI in modular architecture

IPMI v1.5 for Standalone Systems

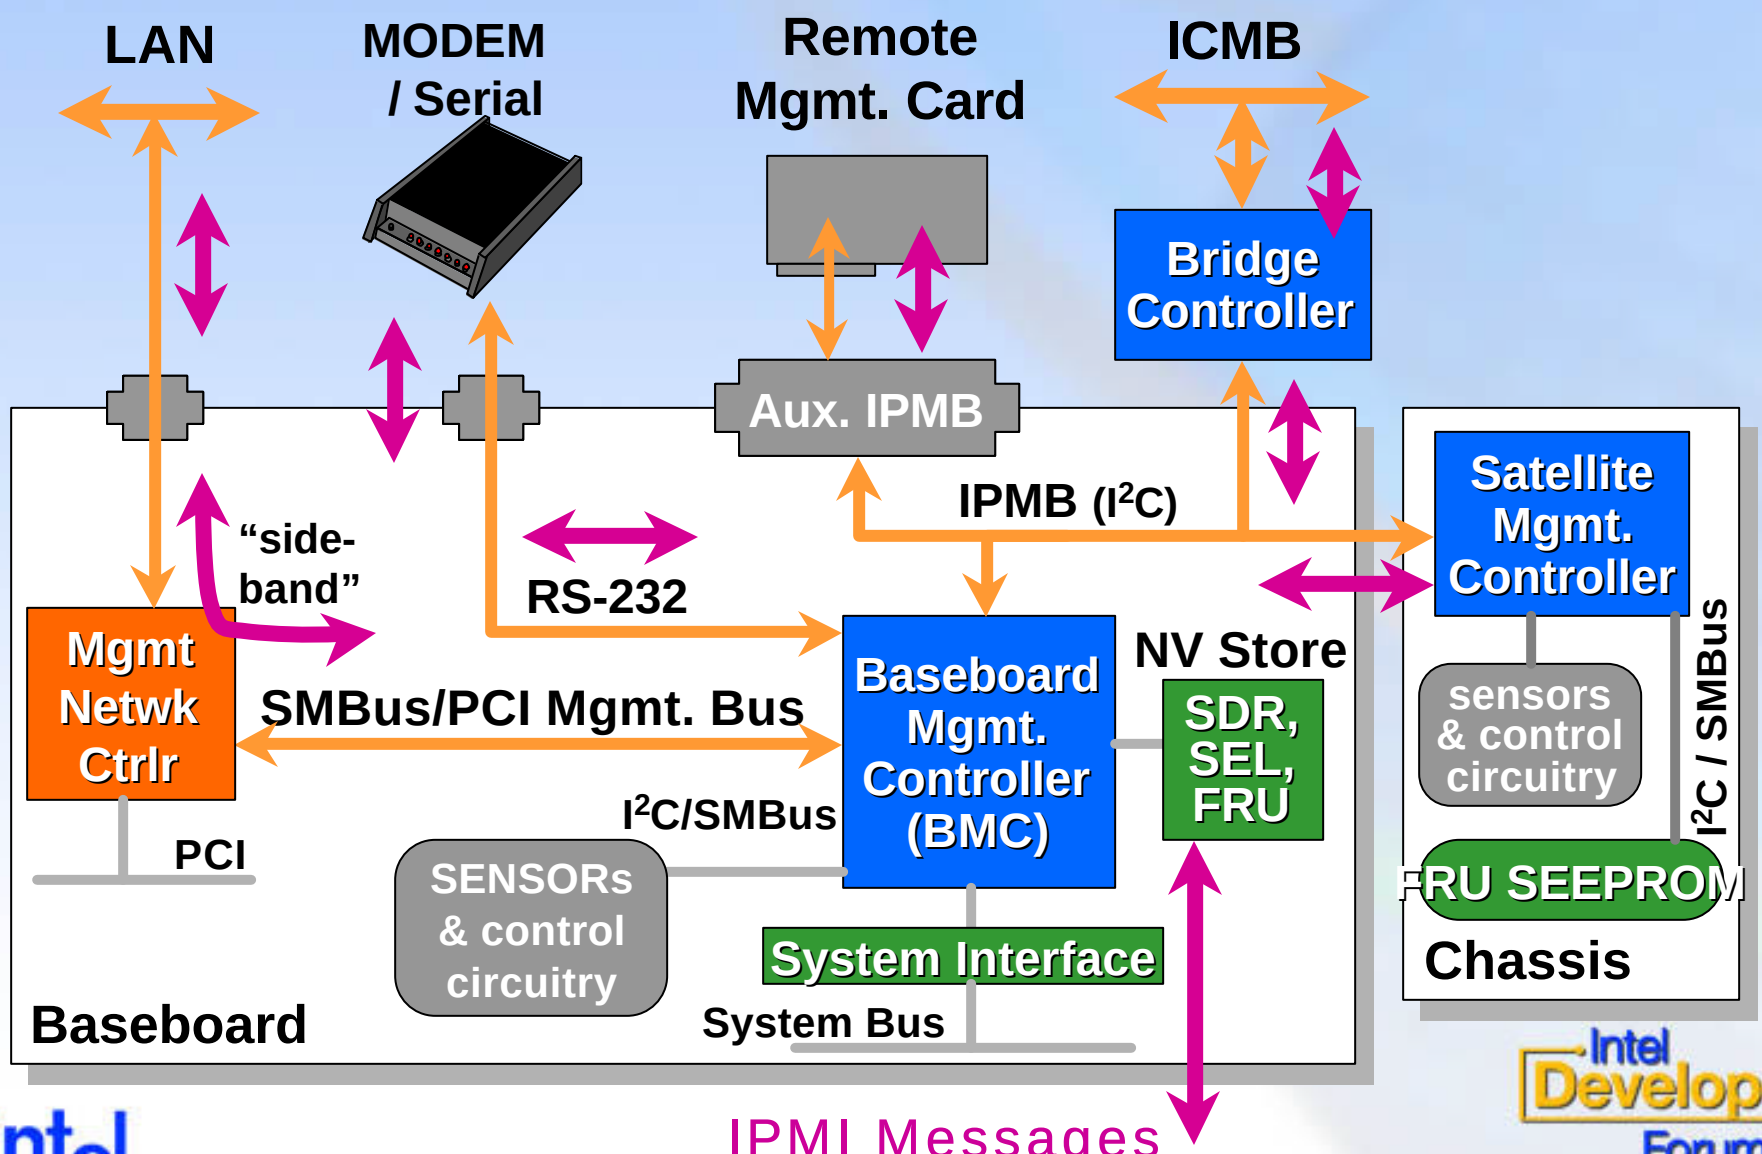

IPMI in modular architecture

Where it fits…

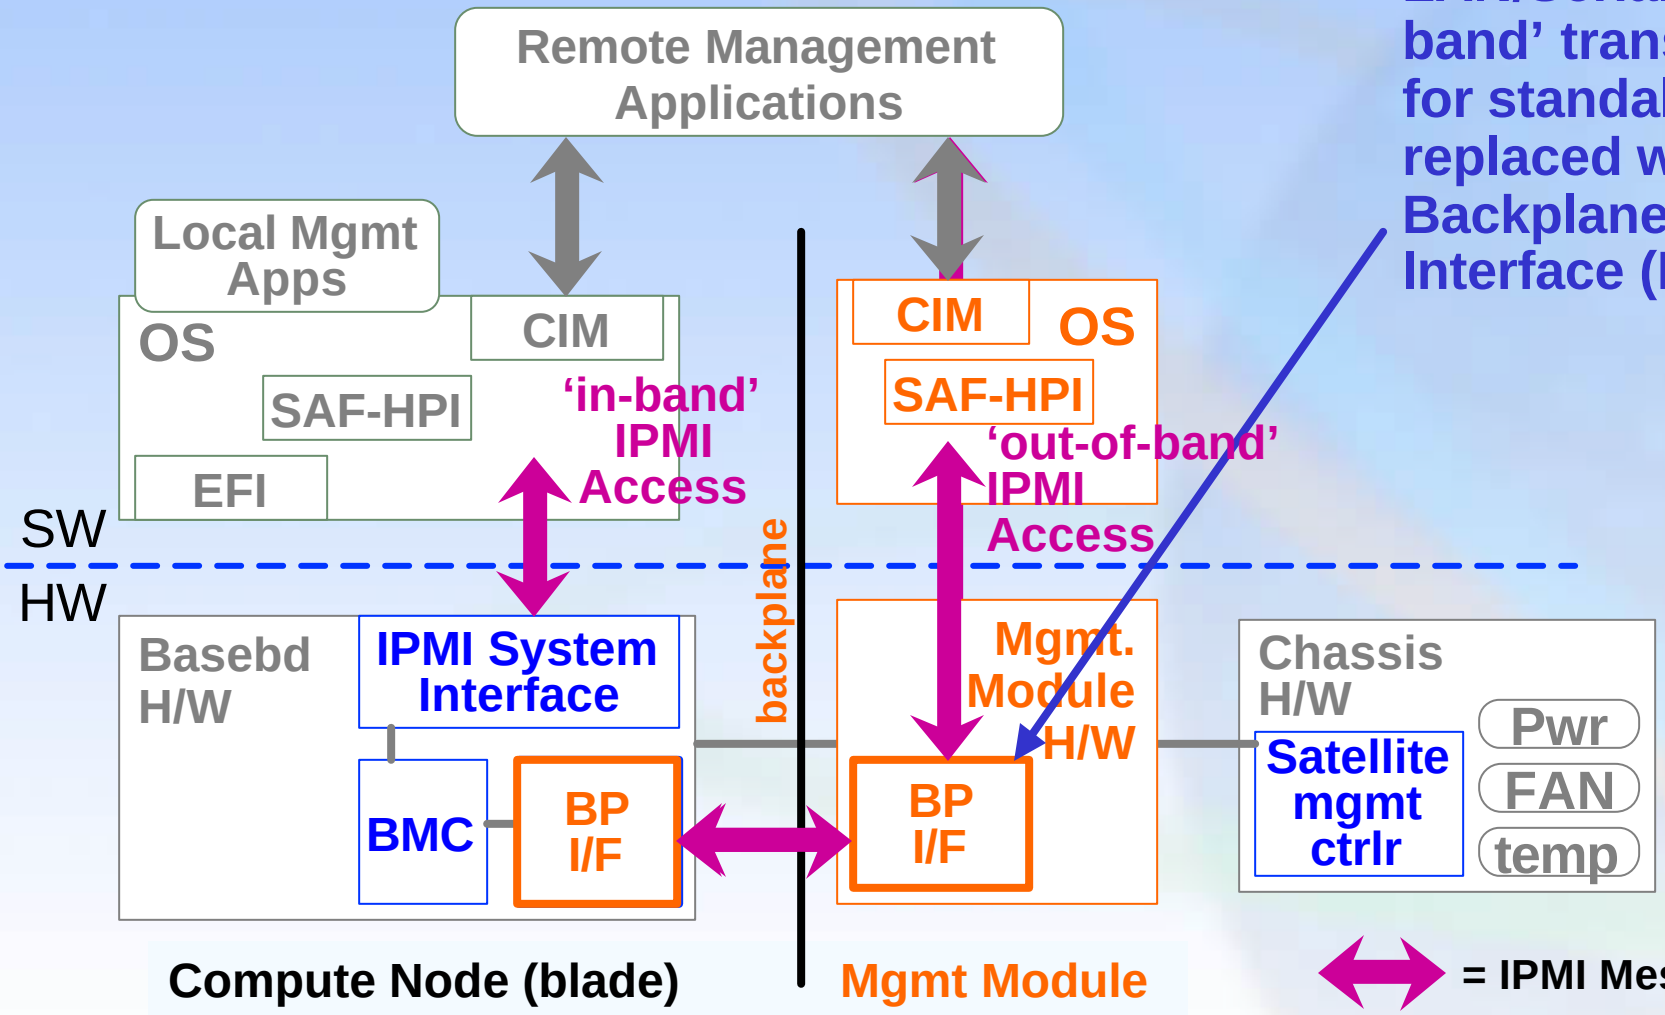

IPMI unifies blade and standalone server management implementations

IPMI in modular architecture

Typical application

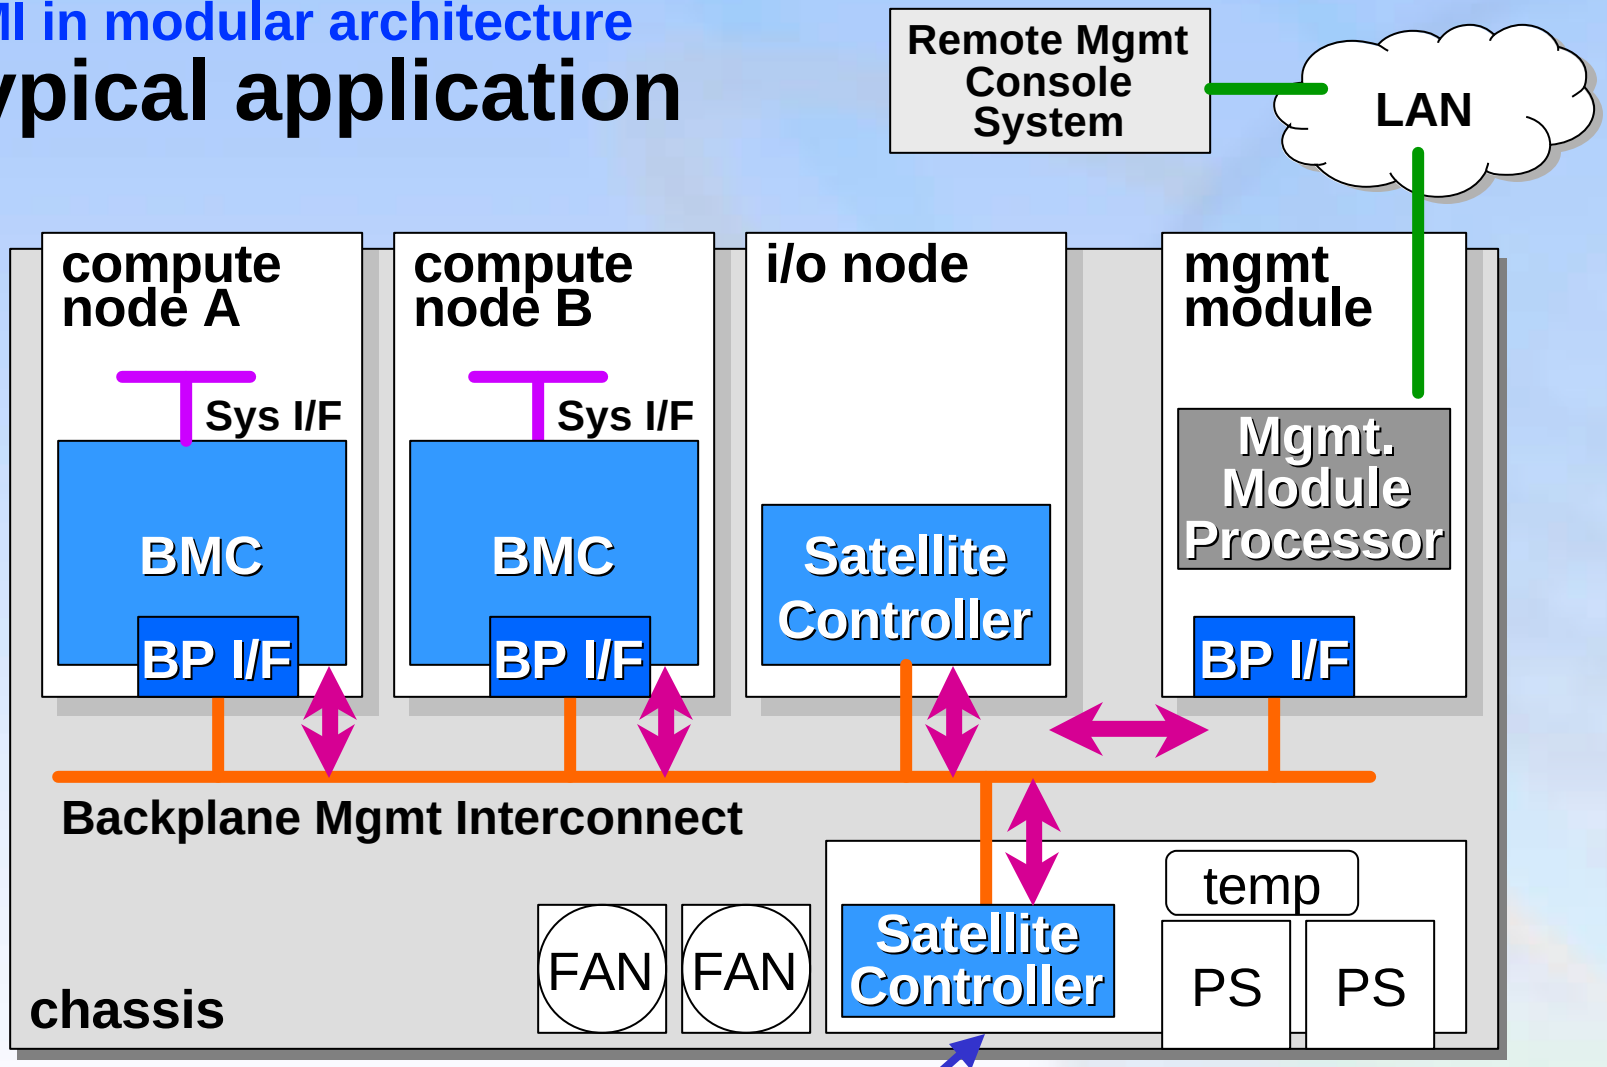

Option: IPMI Satellite Controller

- enables mgmt module re-use across different chassis
- enables any node to be assigned as mgmt module

↔ IPMI Messages

IPMI in modular architecture

Partitioning for protection

- **Problem: Bus topology enables local mgmt s/w to access other nodes**
- **Solution: firmware internal 'firewall'**

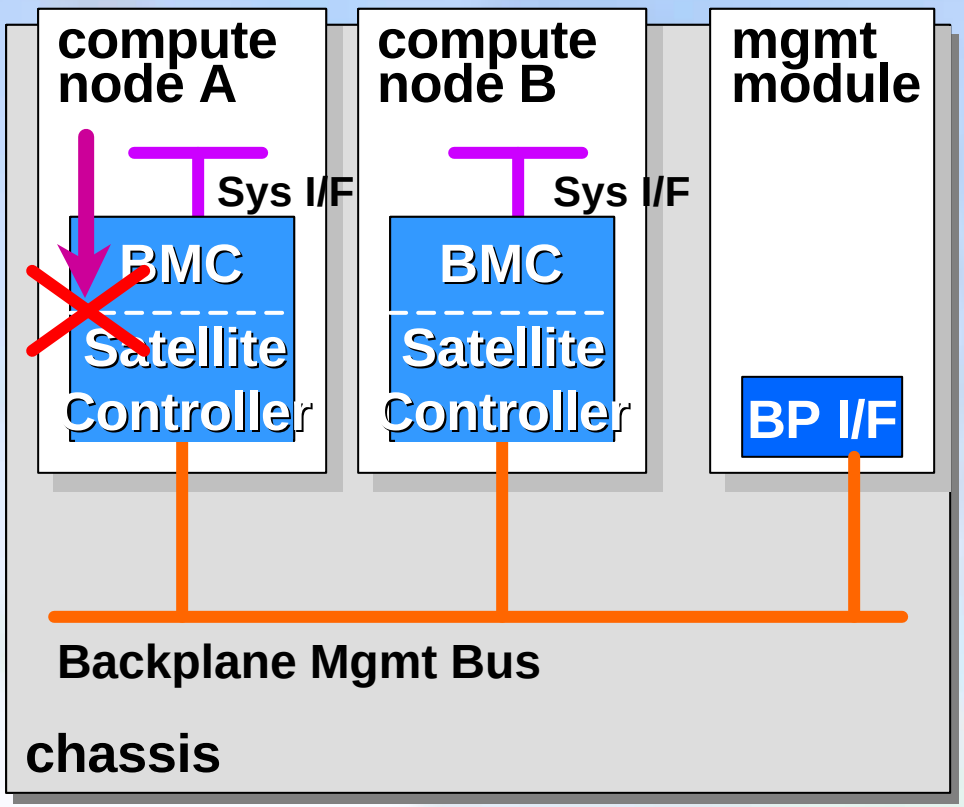

IPMI in modular architecture

Partitioning for protection

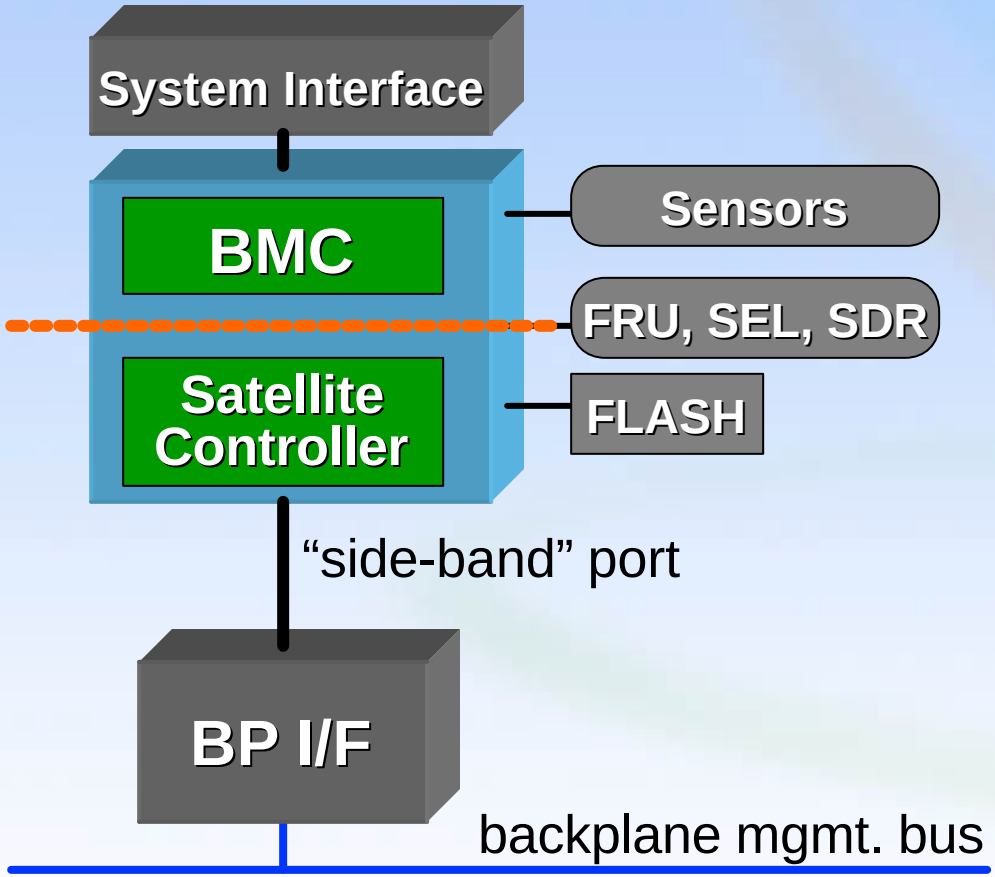

Firmware Internal ‘firewall’

- F/W blocks messaging to other nodes on shared bus
- Allows messages between local software and management module
- Local software may also be blocked from SDR or FRU updates that might be used to generate false events
- Firmware updates can only occur from management bus side
- Access rights can only be configured from management bus side

IPMI in modular architecture

Redundant Backplane Management Busses

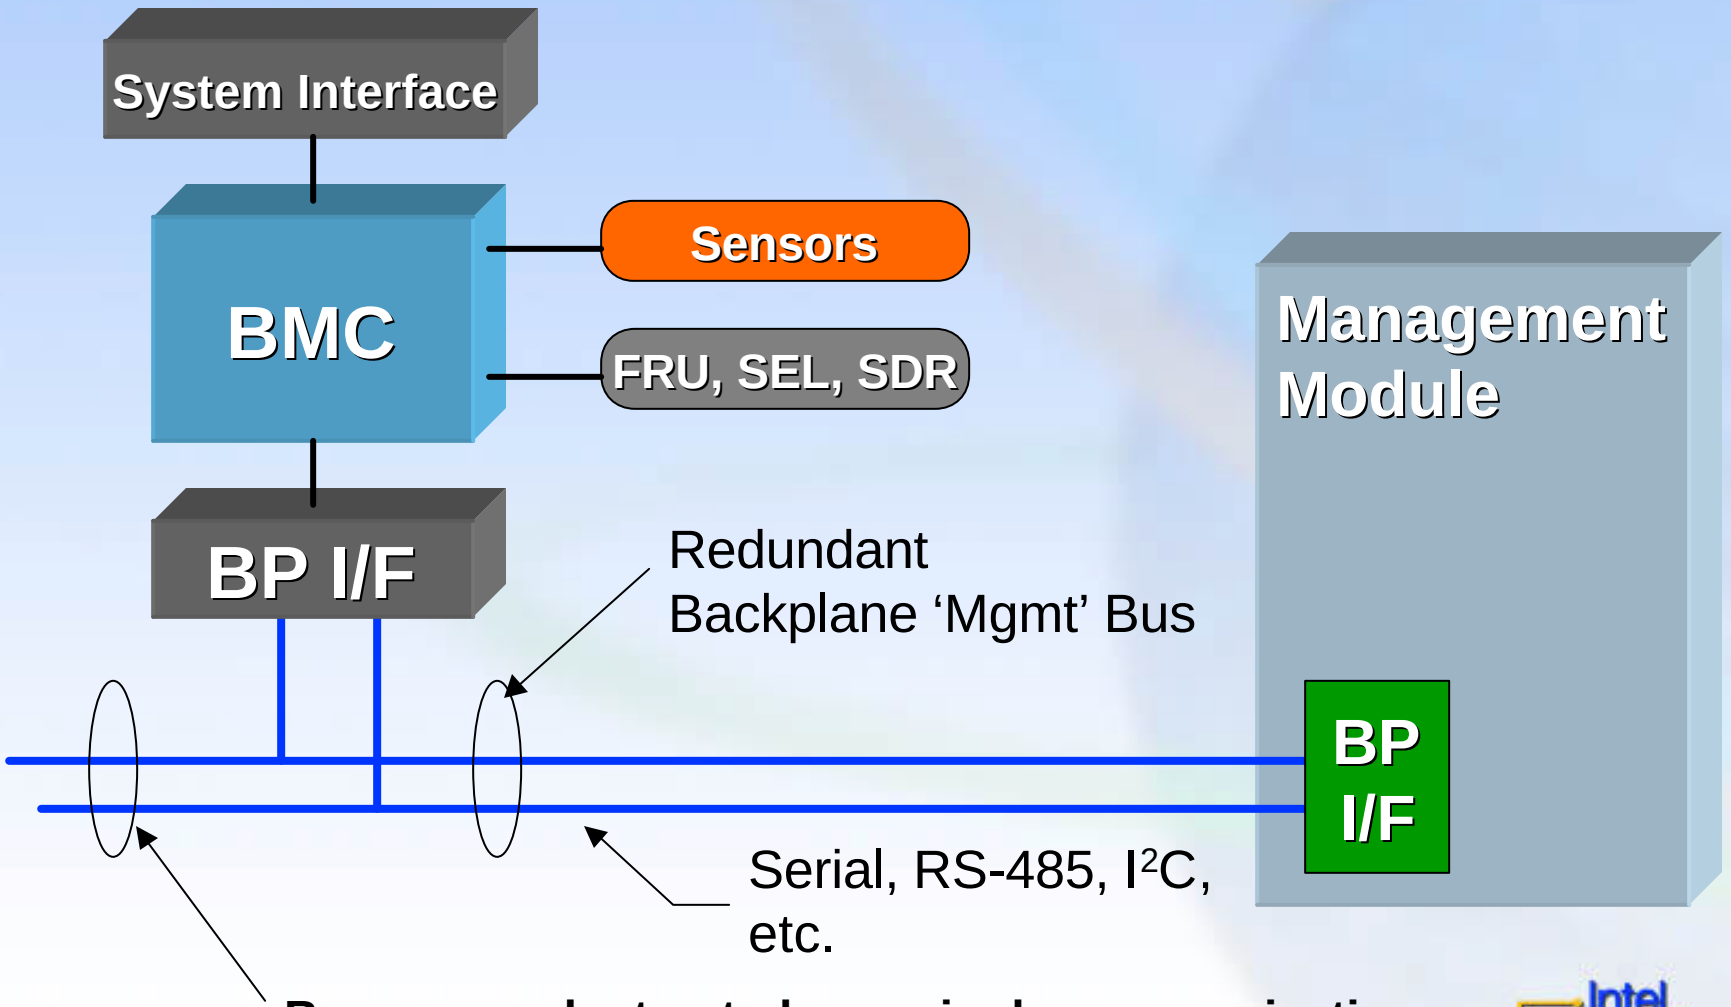

Busses can be treated as a single communication "Channel" under IPMI, or can be separate 'always active' channels

IPMI in modular architecture

Ethernet-based Backplane Mgmt Interface

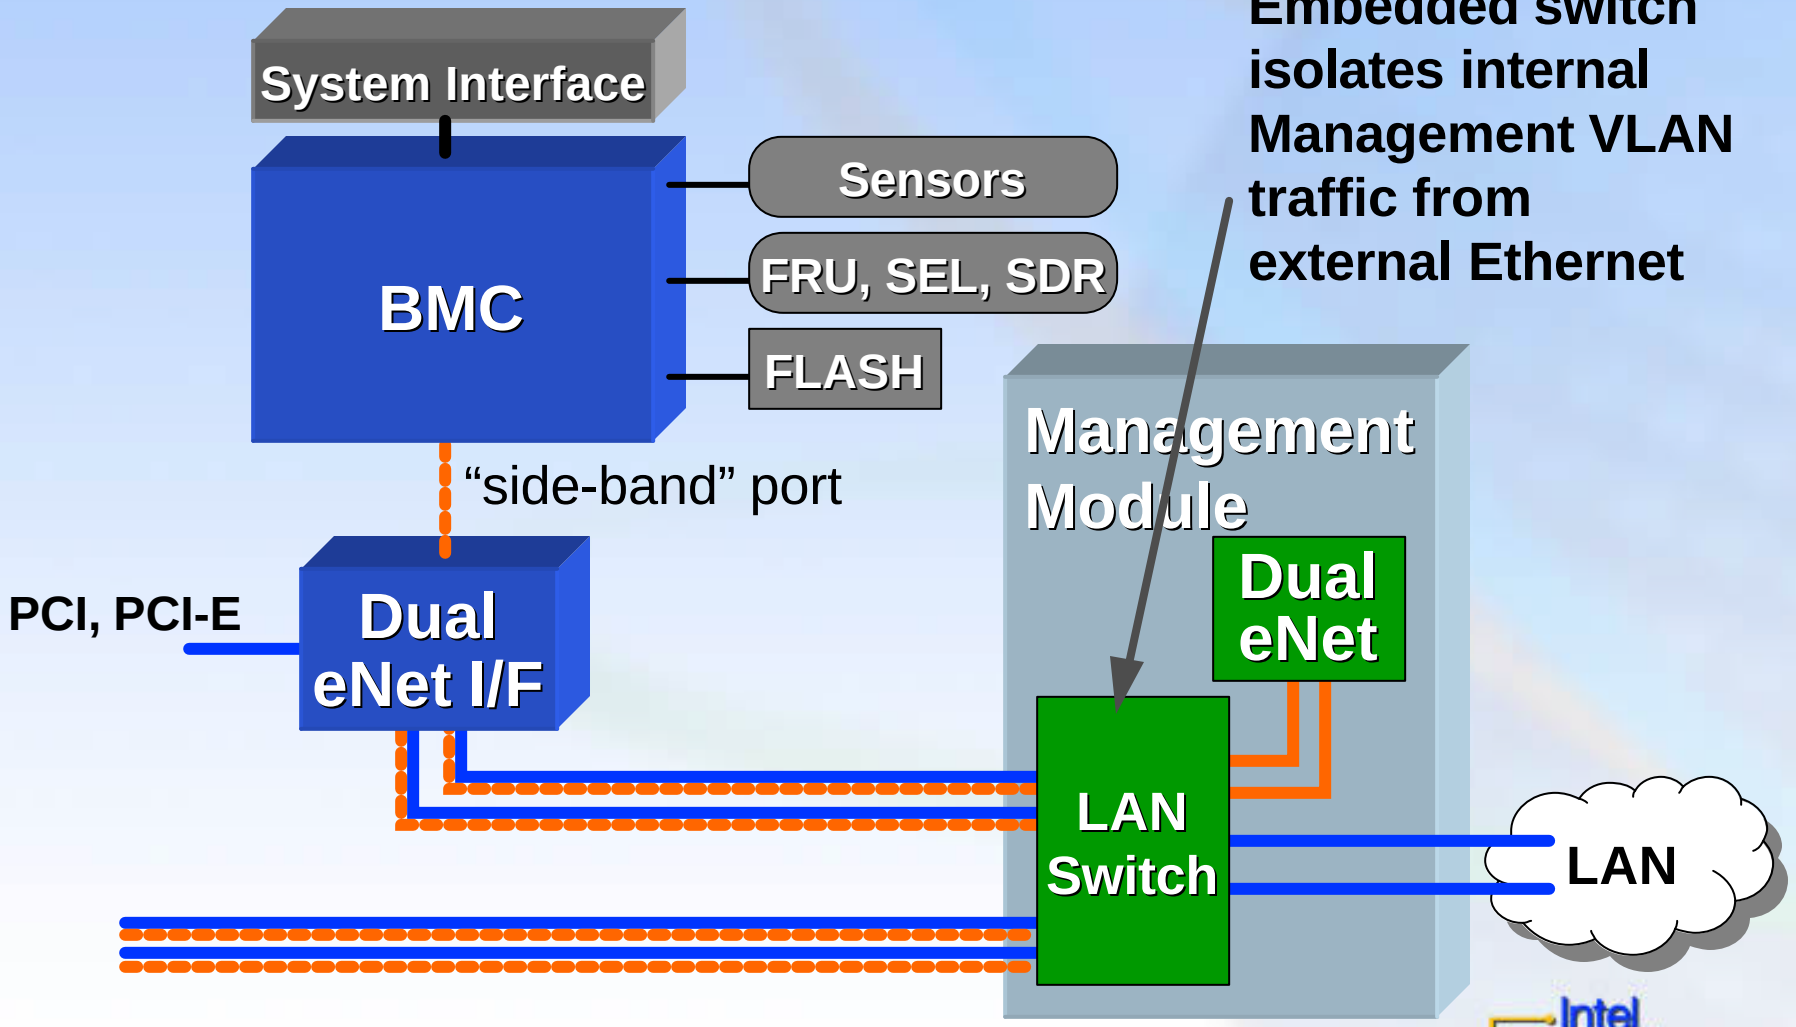

Benefits

- Commonality with standalone servers
  - Reduces ‘special knowledge’ for blade management
- Third party hardware components
  - Management controllers
  - Firmware
- Test tools
  - ICTS: IPMI conformance test suite
    - extensible automated testing for IPMI interfaces
    - can be used for development and validation
- Drivers and Software
  - Linux and Windows operating systems
  - Samples available from IPMI web site
  - Management applications from ISVs

Agenda

- IPMI in modular architecture
- **IPMI specification support for modular**
- IPMI components for modular systems
- IPMI future directions for modular
- Summary

Intel Developer Forum Spring 2003

IPMI Messaging

- **IPMI Channel Model supports multiple paths to BMC**
  - **Supports OEM-defined media**
    - **not limited to present LAN, serial, IPMB**
  - **Supports multiple media types**
    - **e.g. LAN for normal connects, serial/IPMB for backup**
- **Channels can be implemented with redundant physical media**
  - **Channels are a 'logical' concept**
- **Message-bridging architecture provides framework for 'firmware firewall'**
  - **Straightforward parsing to route and filter messages by type**

IPMI support for modular

Node/FRU removal and replacement

Sensor support for coordinating Node and FRU removal and replacement:

- *FRU State* sensor
  - reports hot- or warm- swappable FRU status
    - FRU Not Installed
    - FRU Inactive (in standby or 'hot spare' state)
    - FRU Activation Requested
    - FRU Activation In Progress
    - FRU Active
    - FRU Deactivation Requested
    - FRU Deactivation In Progress
    - FRU Communication Lost
- *Version Change* sensor
  - Reports changes to FRU and/or firmware-software versions or configuration
- *Button/switch* sensor extended for FRU mgmt
  - "FRU Service" and "FRU Latch" offsets

modular system extensions

Under development

- Management Bus ‘failover’ status
  - ability to report redundancy status of an IPMI Channel
- Provisions for ‘Firmware Firewall’
  - spec updates to allow certain functions to be restricted without breaking spec conformance
- Monitoring and control of shared chassis resources: Power, Thermal, Cooling and Slot
  - e.g. support for coordinating chassis FAN speed
  - ‘read only’ access to chassis satellite controller
- Support for node discovery and set up for management
  - e.g. support for reporting ‘node ID’, configuring node access rights to backplane, etc.

Agenda

- IPMI in modular architecture
- IPMI specification support for modular
- **IPMI components for modular systems**
- IPMI future directions for modular
- Summary

Intel Developer Forum Spring 2003

components

Baseboard Management Controllers

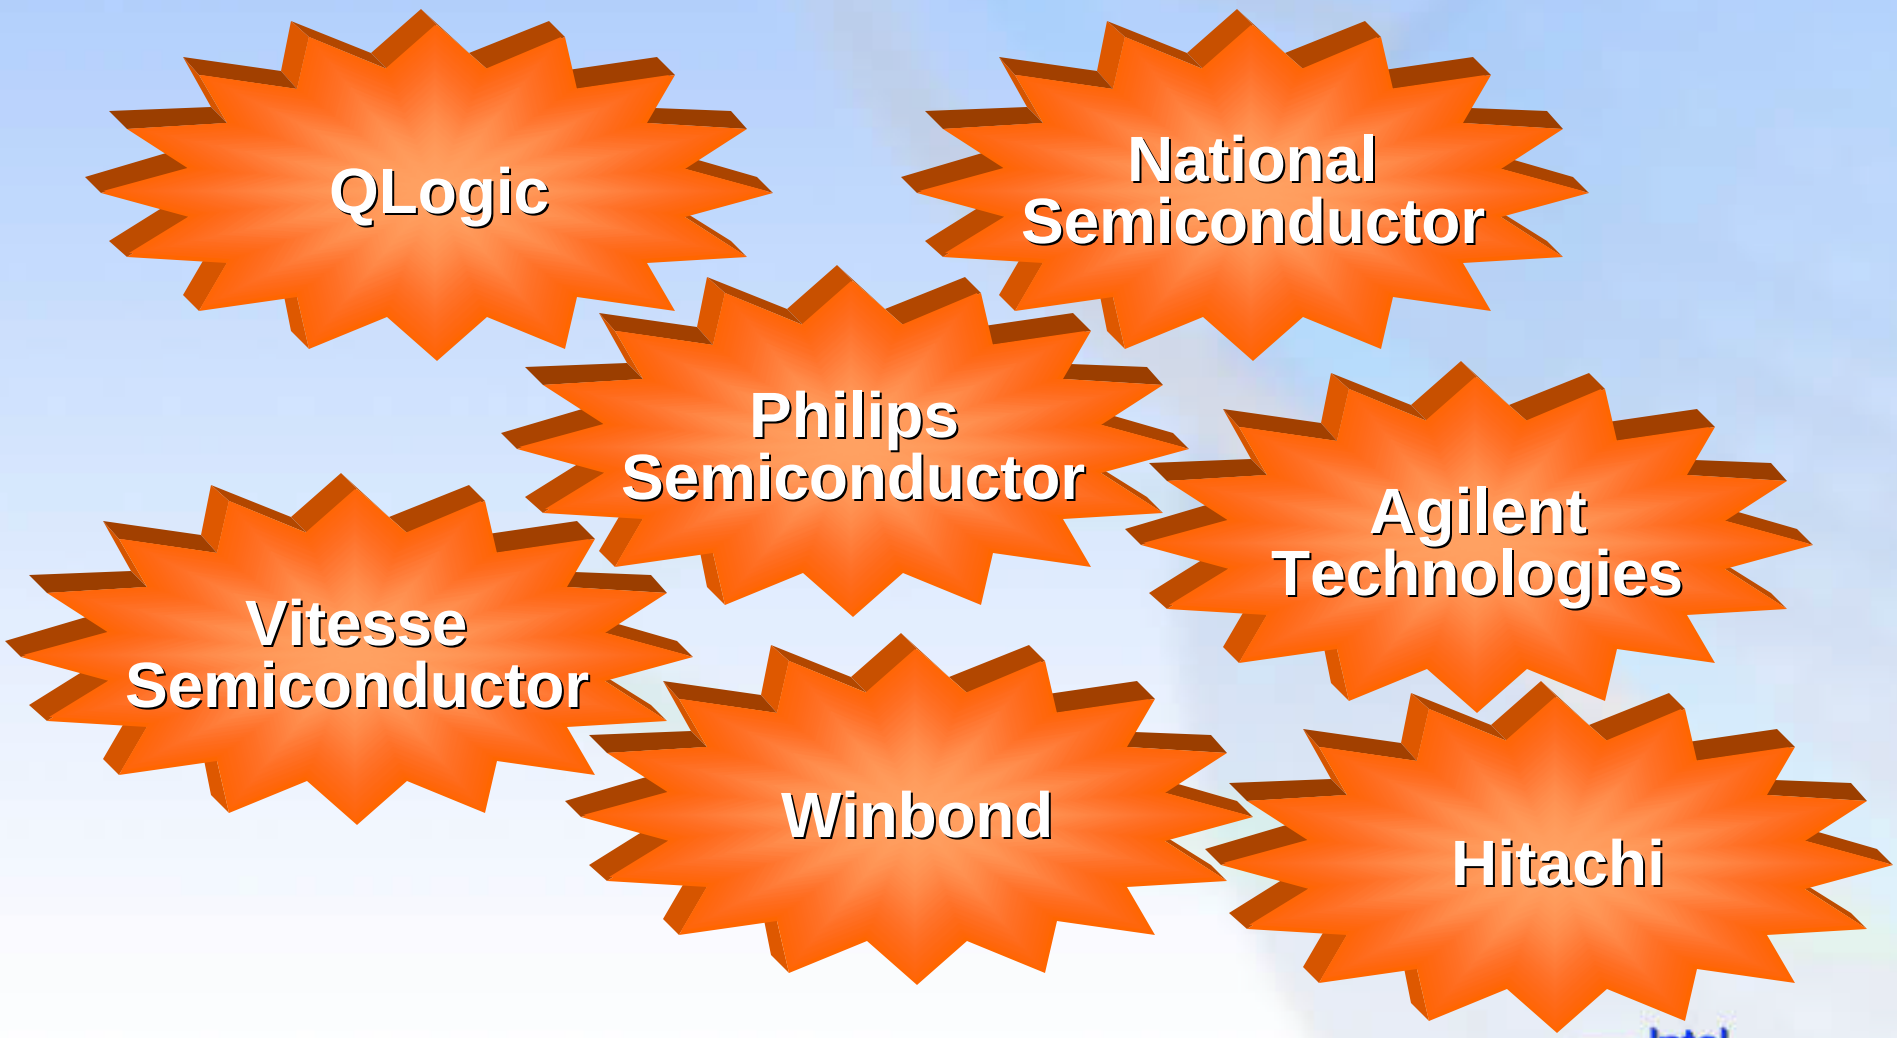

IPMI components fit your design

components

Blade-focused BMCs

Qlogic Zircon UL*

- Target applications: 1P/2P Servers
- 32 GPIO (max), 4 Fan tach, 3 PWM
- 128-pin PQFP
- Virtual Storage Interface – supports ‘virtual floppy’
- Universal Serial Interface
  - for serial redirection / headless
- Production: 4Q02

Qlogic Zircon BL*

- Target applications: server blades, high-end workstations
- 18 GPIO (max), 2 Fan tach, 2 PWM
- 100-pin PQFP
- Virtual Storage Interface
- Universal Serial Interface
- Production: 4Q02

* Other names and brands may be claimed as the property of others.

components

Blade-focused BMCs

- **Hitachi H8S/2145* 'Single-chip' BMC / Satellite Controller**
  - **128KB in-system programmable FLASH and 8KB SRAM on-chip**
  - **Two master-slave 400KHz I2C ports**
  - **Supports 3 KCS LPC channels**
  - **8 A-D, 2 PWM and 4 fan tach inputs, 3 serial ports, up to 75 GP I/O**
  - **Full capability ICE debug is available**
  - **14x14mm TQFP-100 package**
  - **Sampling: Now, Production: May**

components

IPMI Software and Tools

IPMI Web Site

- Reference drivers for Itanium® Architecture and IA-32 under Windows* .NET/2000 and Linux
- IPMI Conformance Test Suite
  - Serves as both validation and development tool

ISV Software

E.g. OSA Technologies

- Management Applications for IPMI
  - "Remote Console" applications and IPMI Drivers
- Firmware engineering also available
  - SDKs for popular BMCs
  - Supports IPMI v1.5 and out-of-band access (serial, LAN)

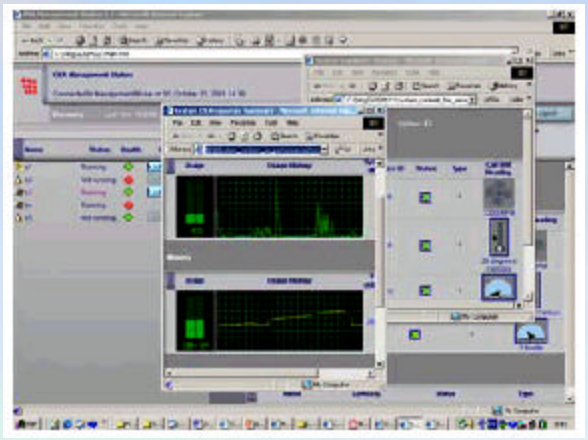

IPMI components improve TTM and reduce design cost

* Other names and brands may be claimed as the property of others.

Agenda

- IPMI in modular architecture
- IPMI specification support for modular
- IPMI components for modular systems
- **IPMI future directions for modular**
- Summary

IPMI Future Directions

Technology Transitions

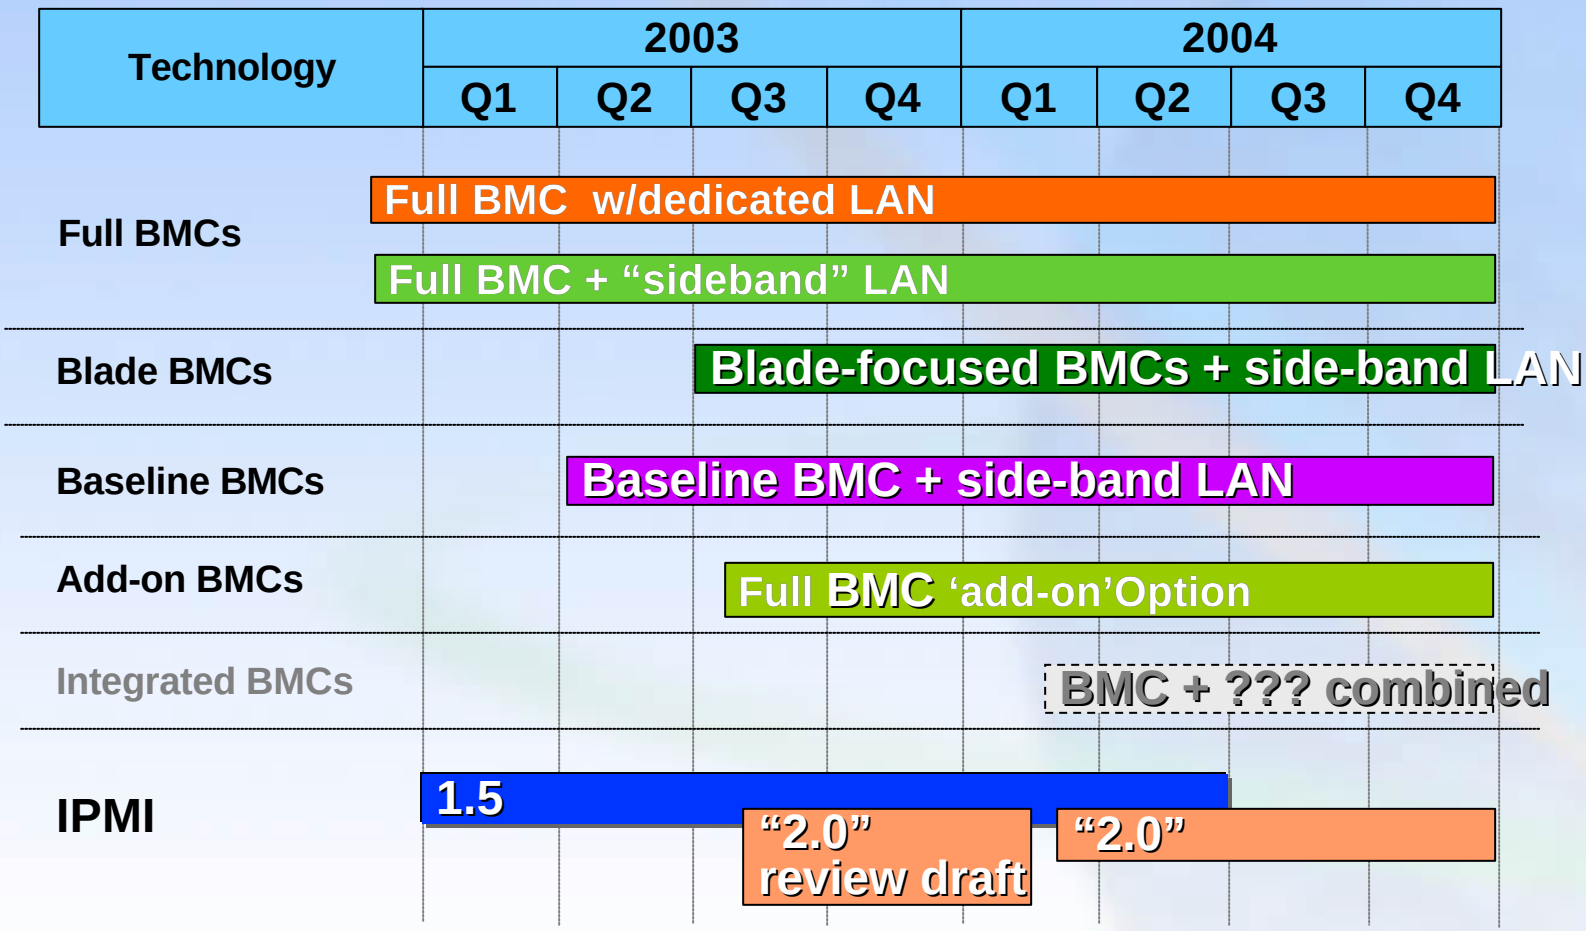

Update next IDF

IPMI Future Directions

IPMI '2.0' Proposals

- Serial redirection over LAN
- Terminal mode extensions (improved 'CLI')
- ASF Alignment
  - Common authentication protocols
  - Smoothes ASF to IPMI transition between desktop and sub-entry server systems
- Modular (blade) support
  - blade/chassis relationship, blade power mgmt., etc.
  - AdvancedTCA support (formerly 'CompactPCI')
- IPMI over Web
  - enabling technology for IPMI over Web (may be post 2.0)

IPMI continues to evolve valuable new capabilities

Agenda

- IPMI in modular architecture
- IPMI specification support for modular
- IPMI components for modular systems
- IPMI future directions for modular
- **Summary**

Intel Developer Forum Spring 2003

Summary

- IPMI unifies modular and general-purpose server platform management
- IPMI components and tools save design time and improve TTM
- IPMI provides the foundation for future platform management initiatives and features

Acronyms

- 3GIO – Third Generation IO, now named PCIExpress*
- 10GbE – 10 Gigabit Ethernet
- ACPC - Automatic Control of Power Consumption
- ASF – Alert Standard Format http://www.dmtf.org/standards/standard_alert.php
- aTCA - Advanced Telecom Computing Architecture*
- B2B – Business to Business
- BE – Enterprise Backend (Data Services)
- BIS – Boot Integrity Service http://www.intel.com/design/security/bis/biswks.htm
- BMC – Baseboard Management controlller
- CERT – Certificate like X.509 http://www.ietf.org/html.charters/pkix-charter.html
- CLR – common language runtime - engine running MSIL
- COO – Cost of Ownership
- cPCI – CompactPCI* (PICMG 2.x)
- DBS – Demand Based Switching
- EAP – Extensible Authentication Protocol http://www.faqs.org/rfcs/rfc2284.html
- EFI – Extensible Firmware Interface http://www.intel.com/technology/efi/main_specification.htm
- EMCI – Enterprise Modular Computing Initiative
- EPTM - Enterprise Power and Thermal Manager FE – Front End
- FRU – Field Replaceable Unit
- GbE – Gigabit Ethernet
- ICMB - Intelligent Chassis Management Bus
- IPMB -Intelligent Platform Management Bus
- IPMI – Intelligent Platform Management Interface http://www.intel.com/design/servers/ipmi/index.htm
- ISA – Instruction Set Architecture
- ISCCI – Internet SCSI (Small Computer System Interface)
- J2EE – Java 2 Enterprise Edition
- JVM – Java Virtual Machine

Acronyms

- L3 – level 3
- MT – Enterprise Mid-Tier (Application Services)
- MP – multiprocessing (4P & above)
- MRTE – Managed Run-Time Environments
- MSIL – Microsoft intermediate language; compiler language of .Net code
- OOB – Out of Band
- PCI - Peripheral Component Interconnect
- PEF - Platform Event Filtering
- PICMG3.x- PCI Industrial Computer Manufacturer's Group - Follow on to cPCI (PICMG 2.x)
- PXE – Preboot eXecution Environment http://www.intel.com/technology/efi/main_specification.htm
- RADIUS – Remote Authentication Dial-In User Service http://www.faqs.org/rfcs/rfc2138.html
- RCMP – Remote Management and Control Protocol http://www.dmtf.org/standards/standard_alert.php
- RCO – Real Cost of Ownership
- SAF - Service Availability Forum
- SAF-HPI - Server Availability Forum - Hardware Platform Interface
- SAF-TE - SCSI Accessed Fault-Tolerant Enclosures.
- SDR - Sensor Data Record
- SEL - System Event Log
- TCO – total cost of ownership
- TOE – TCP/IP offload engine
- TDP – thermal design point
- UD – Ultra Dense
- U – unit of measure of server height (1U = 1.75 inches)
- WBEM – Web-Based Enterprise Management http://www.dmtf.org/standards/standard_wbem.php
- WfM – Wired for Management http://www.intel.com/labs/manage/wfm/index.htm
- 2P – dual processing capable
- 4P – quad processing capable

Collateral

- White paper: Modular Computing: The New Enterprise Computing Model (Egenera/Intel)
- URLS:
  - IPMI Web Site: http://developer.intel.com/design/servers/ipmi
  - Distributed Management Task Force (DMTF): http://www.dmtf.org
  - IBM Autonomic Computing*: http://www-3.ibm.com/autonomic/index.shtml
  - IBM eLiza* project on X-series: http://www-1.ibm.com/servers/autonomic/
  - IBM BladeCenter*: http://www.pc.ibm.com/us/eserver/xseries/bladecenter_family.html?ca=xSeries&met=ibmblade&me=A
  - HP Utility Computing*: http://devresource.hp.com/topics/utility_comp.html
  - Microsoft .NET*: http://www.microsoft.com/net/
  - Egenera*: http://www.egenera.com/prod_spec_valprop.php
  - Sun N1*: http://wwws.sun.com/software/solutions/n1/index.html
  - Giga* analyses: (R.Fichera)
    - Criteria for Selection: Bladed and Modular Servers (July 31, 2002)
    - Future of the Data Center: Modularity and Virtualization (May 8, 2002)
    - Economics of Cable Consolidation: A Major Impact on Server Cost (July 23, 2002)

Glossary

BMC	Baseboard Management Controller.
FRU	Field Replaceable Unit. A field replaceable component such as a board, module, fan, power supply, etc.
ICMB	Intelligent Chassis Management Bus. The ICMB provides a dedicated management bus that enables delivering IPMI messages and alerts between multiple host and peripheral chassis.
IPMB	Intelligent Platform Management Bus. Name for the architecture, protocol, and implementation of a special bus that interconnects the baseboard and chassis electronics and provides a communications media for system platform management information.
IPMI	Intelligent Platform Management Interface. IPMI defines a common, abstracted, and self-descriptive interface for platform management hardware that monitors server characteristics such as temperature, voltage, fans, power supplies, and chassis.
OOB	Out-of-Band. System platform management access that does not involve going through the OS or other software running on the main processors of the managed system.
PEF	Platform Event Filtering. A feature in IPMI that enables the BMC to generate a selectable action (e.g. power on/off, reset, send Alert, etc.) when a configurable event occurs on the management system.

Glossary

SAF	Service Availability Forum. Standards body consisting of Telco platform and software vendors that is defining RAS standards including UCMI
SAF-HPI	Server Availability Forum - Hardware Platform Interface. Name for a set of APIs and structures for representing and accessing platform management hardware.
SAF-TE	SCSI Accessed Fault-Tolerant Enclosures. SAF-TE provides a mechanism that enables RAID fault information to be sent to the hot-swap backplane via SCSI.
SDR	Sensor Data Record. SDRs provide the information that tells management software what sensors, events, management controllers, and FRU information is available from a given IPMI implementation.
SEL	System Event Log. A non-volatile storage area and associated interfaces for storing system platform event information for later retrieval.


Intel
Developer
Forum
Spring 2003